Exercise 5

1 Epidemic Probabilistic Broadcast

Consider the system model of Section 5.5, where up to t of n servers may crash. When n is large (thousands or even millions), protocols of the kind we have considered perform poorly, because a single server needs to send messages to and receive messages from all others.

An alternative is offered by hierarchical protocols, where messages are not sent to all servers in the system at once, but disseminated by a set of relaying servers organized in a tree. Because of faults, however, one would either have to construct a fault-tolerant tree or use complicated tree adjustment method based on failure detectors.

As a simpler alternative, *epidemic* protocols for reliable broadcast have been proposed, where the sender sends its message to a small randomly chosen subset of αn servers, who in turn forward the message to another randomly chosen set of αn, etc. This process continues for a number of k rounds.

Since there will be runs where not all servers deliver the message, the *validity* and *agreement* conditions of reliable broadcast have to be weakened to the following:

Probabilistic validity: There is a constant ϵ such that if a correct sender *broadcasts* a message m, then every correct servers *delivers* m with probability at least $1 - \epsilon$.

Tasks:

1. Formalize a protocol for epidemic reliable broadcast.

2. Derive a lower bound on the probability that s servers have delivered the message after k rounds. (Use Maple or Mathematica to plot it.)